

CLAIMS

What is claimed is:

1. A method for rate-based flow control between a sender and a receiver,
the method comprising:
 - 5 (a) at a sender, sending packetized information to a receiver at a first rate;
 - (b) at the receiver:
 - (i) receiving the packetized information;
 - (ii) computing a congestion window size and a round-trip time
10 based on the packetized information;
 - (iii) computing a transmission rate based on the congestion window size and the round-trip time;
 - (iv) periodically transmitting the transmission rate to the sender;
and
 - 15 (c) at the sender, controlling the rate for sending the packetized information to the receiver based on the transmission rate received from the receiver.
2. The method of claim 1 wherein computing a congestion window size
comprises computing an average congestion window size over a
20 predetermined time interval, computing a round-trip time comprises
computing an average round-trip time over the time interval, and
computing the transmission rate includes dividing the average congestion
window size by the average round-trip time.
3. The method of claim 1 wherein computing a congestion window size at
25 the receiver includes incrementing the congestion window size by one

segment in response to receiving a properly-sequenced packet from the sender.

4. The method of claim 1 wherein computing a congestion window size at the receiver includes maintaining a current congestion window size in response to receiving an improperly sequenced packet from the sender.

5. The method of claim 1 comprising, at the receiver, implementing a state machine including TCP-sender-like flow control states for adjusting the congestion window size.

6. The method of claim 5 wherein implementing a state machine includes implementing a state machine having a slow start state in which the congestion window size is incremented by one segment in response to receiving a properly-sequenced packet from the sender.

7. The method of claim 6 wherein implementing a state machine includes implementing a state machine having a congestion avoidance state in which the congestion window size is increased by the inverse of a previous congestion window size in response to receiving a properly-sequenced packet from the sender.

8. The method of claim 7 wherein implementing a state machine comprises implementing a state machine including a gap state reachable from the slow start state and the congestion avoidance state in response to receiving an improperly sequenced packet from the sender.

9. The method of claim 8 comprising, in response to receiving a packet that triggered transition to the gap state, transitioning to the state that the receiver was in prior to entering the gap state.

10. The method of claim 8 wherein implementing a state machine comprises

implementing a state machine having a fast recovery state reachable from the gap state in which the receiver reduces the congestion window size only once in response to multiple packet losses within a single congestion window.

- 5 11. The method of claim 1 wherein the sender adjusts the transmission rate without receiving per-packet acknowledgements from the receiver.
12. The method of claim 8 wherein computing a congestion window size comprises computing an average congestion window size over a predetermined time interval, computing a round-trip time comprises
10 computing an average round-trip time over the time interval, and computing the transmission rate includes dividing the average congestion window size by the average round-trip time.
13. The method of claim 12 comprising dynamically adjusting the predetermined time period based on the state of the receiver.
- 15 14. The method of claim 13 wherein dynamically adjusting the predetermined time period includes setting the time period equal to a value equal to the time difference between a first time when the receiver enters the slow start or congestion avoidance state and a second time when the receiver re-enters the slow-start or congestion avoidance state.
- 20 15. The method of claim 1 wherein computing a roundtrip time includes estimating the roundtrip time based on the time to receive packets from the sender equal to a current congestion window size.
16. The method of claim 1 wherein computing the transmission rate includes computing the transmission rate based on a weighted average of a
25 plurality of congestion window sizes divided by corresponding roundtrip

times.

17. The method of claim 16 wherein computing a weighted average includes weighting recent congestion window sizes more heavily than older congestion window sizes.

5 18. A receiver-based system for controlling flow of packetized data between a sender and a receiver, the system comprising:

(a) a sender for sending packetized data over a network; and

(b) a receiver for receiving the packetized data from the sender and for emulating TCP-sender-based flow control algorithms to compute a transmission rate and for forwarding the transmission rate to the sender, wherein the sender adjusts its rate for sending the packetized data to the receiver based on the transmission rate.

10 19. The system of claim 18 wherein the receiver is adapted to compute a congestion window size and a round-trip time based on packets received from the sender.

15 20. The system of claim 19 wherein the receiver is adapted to compute an average congestion window size and an average round-trip time over a predetermined time interval and to compute the transmission rate based on the average congestion window size and the average round-trip time, thereby smoothing fluctuations in the transmission rate.

20 21. The system of claim 19 wherein the receiver is adapted to implement a state machine including TCP-sender-like flow control states for adjusting a congestion window size used to compute the transmission rate.

25 22. The system of claim 21 wherein the state machine includes a slow start state in which the congestion window size is incremented by one segment

in response to receiving a properly-sequenced packet from the sender.

23. The system of claim 22 wherein the state machine includes a congestion avoidance state in which the congestion window size is increased by the inverse of a previous congestion window size in response to receiving a properly-sequenced packet from the sender.

24. The system of claim 23 wherein the state machine includes a gap state reachable from the slow start state and the congestion avoidance state in response to receiving an improperly sequenced packet from the sender.

25. The system of claim 24 wherein, in response to receiving a packet that triggered transition to the gap state, the receiver is adapted to transition to the state that the receiver was in prior to entering the gap state.

26. The system of claim 24 wherein the state machine includes a fast recovery state reachable from the gap state in which the receiver reduces the congestion window size only once in response to multiple packet losses within a single congestion window.

27. The system of claim 18 wherein the sender is adapted to adjust its rate for sending packetized data to the receiver without receiving per-packet acknowledgements from the receiver.

28. The system of claim 24 wherein the receiver is adapted to compute an average congestion window size and an average round-trip time over a predetermined time interval and to compute the transmission rate based on the average congestion window size and the average round-trip time, thereby smoothing fluctuations in the transmission rate.

29. The system of claim 28 wherein the receiver is adapted to dynamically adjust the predetermined time period based on the state of the receiver.

30. The system of claim 29 wherein dynamically adjusting the predetermined time period includes setting the time period equal to a value equal to the time difference between a first time when the receiver enters the slow start or congestion avoidance state and a second time when the receiver re-enters the slow-start or congestion avoidance state.
- 5
31. The system of claim 15 wherein the receiver is adapted to compute the transmission rate based on a weighted average of congestion window sizes.